arXiv:1402.6035v1 [stat.ME] 25 Feb 2014

Annealed Important Sampling for Models with Latent Variables

M.-N. Tran* C. Strickland* M. K. Pitt† R. Kohn*

February 26, 2014

Abstract

This paper is concerned with Bayesian inference when the likelihood is analytically intractable but can be unbiasedly estimated. We propose an annealed importance sampling procedure for estimating expectations with respect to the posterior. The proposed algorithm is useful in cases where finding a good proposal density is challenging, and when estimates of the marginal likelihood are required. The effect of likelihood estimation is investigated, and the results provide guidelines on how to set up the precision of the likelihood estimation in order to optimally implement the procedure. The methodological results are empirically demonstrated in several simulated and real data examples.

Keywords. Intractable likelihood; Latent variables; Sequential Monte Carlo; Unbiasedness; Marginal likelihood

*Australian School of Business, University of New South Wales
†Department of Economics, University of Warwick

1 Introduction

Models incorporating latent variables are very popular in many statistical applications. For example, generalized linear mixed models (Fitzmaurice et al., 2011), that use latent variables to account for dependence between observations, appear in the genetics, social and medical sciences literatures as well as many other areas of statistics. State space models (Durbin and Koopman, 2001), whose latent variables follow a Markov process, are used in economics, finance and engineering. Gaussian process classifiers (see, e.g., Filippone, 2013; Rasmussen and Williams, 2006), that use a set of latent variables distributed as a Gaussian process to account for uncertainty in predictions, are used in computer science.

Inference about the model parameters θ in latent variable models can be challenging because the likelihood is expressed as an integral over the latent variables. This integral is analytically intractable in general. It can also be computationally challenging when the dimension of the latent variables is high. Recent work by Beaumont (2003); Andrieu and Roberts (2009); Andrieu et al. (2010) shows that it is possible to carry out Bayesian inference in latent variable models by using an unbiased estimate of the likelihood within a Markov chain Monte Carlo (MCMC) sampling scheme. This method is known as particle MCMC. However, the resulting Markov chain can often be trapped in local modes and it is difficult to asses if it has converged. We find that even for simple models it is difficult for the Markov chain to mix adequately and the chain can take a long time to converge if the log of the estimated likelihood is too variable. The marginal likelihood is often used to choose between models. Another drawback of MCMC in general, and particle MCMC in particular, is that it is often difficult to use it to estimate the marginal likelihood.

Another approach to Bayesian inference for models with latent variables is importance sampling squared (IS^2) proposed in Tran et al. (2013). They show that importance sampling (IS) with the likelihood replaced by its unbiased estimate is still valid for estimating expectations with respect to the posterior. IS^2 offers several advantages over particle MCMC: (1) It is easy to estimate the standard errors of the estimators; (2) It is straightforward to parallelize

the computation; (3) It is straightforward to estimate the marginal likelihood. However, as is typical of importance sampling algorithms, a potential drawback with IS^2 is that its performance may depend heavily on the proposal density for θ. A good proposal density may be difficult to obtain in complex models.

When it is possible to evaluate the likelihood, annealed importance sampling (AIS) (Neal, 2001) is a useful method for estimating expectations with respect to the posterior for θ. AIS is an importance sampling method in which samples are first drawn from an easily-generated distribution and then moved towards the distribution of interest through Markov kernels. AIS explores the parameter space efficiently and is useful in cases where the target distribution is multimodal and/or when choosing an appropriate proposal density is challenging.

This article proposes an AIS algorithm for Bayesian inference when working with an estimated likelihood, which we denote as AISEL (annealed importance sampling with an estimated likelihood). Our first contribution is to show that the algorithm is valid for estimating expectations with respect to the exact posterior when the likelihood is estimated unbiasedly. As with particle MCMC and IS^2, it is important to understand the effect of estimating the likelihood on the resulting inference. The second contribution of this article is to answer this question by comparing the efficiency of AISEL with the efficiency of the corresponding AIS procedure that assumes that the likelihood is available. We show that the ratio of the efficiency of AISEL to that of AIS is smaller than or equal to 1, and the ratio is equal to 1 if and only if the estimate of the likelihood is exact. This ratio decreases exponentially with the product of the variance of the log of the estimated likelihood and a term that depends on the the annealing schedule in the AIS algorithm. The term based on the annealing schedule is small if the annealing schedule evolves slowly. This result allows us to understand how much accuracy is lost when working with an estimated likelihood. An attractive feature of AISEL is that it is more robust than IS^2 and particle MCMC to the variability of the log likelihood estimate. This is important when only highly variable estimates of the likelihood are available, which often occurs if it is expensive to obtain accurate estimates of the likelihood.

The third contribution of the article is to provide theory and practical guidelines for

optimally choosing the number of particles to estimate the likelihood so as to minimize the overall computational cost for a given precision. The fourth contribution is to describe an efficient yet simple method to compute the marginal likelihood, which is important for model choice.

The SMC^2 algorithm of Chopin et al. (2013) sequentially updates the posterior of the model parameters as new observations arrive. The validity of the method is justified because the likelihood estimated by the particle filter is unbiased. In contrast, our AIS algorithm uses all the data and is static. As discussed in their paper (Chopin et al., 2013, Section 5.2), the reasoning used to justify SMC^2 does not apply to tempered sampling in the spirit of the AIS, because a tempered likelihood estimator is not an unbiased estimator of the corresponding tempered likelihood. However, Section 3.1 uses variable augmentation to justify the validity of the AIS method when working with an unbiased likelihood estimate.

We illustrate the proposed methodology through a simulated example, as well as the analysis of a Pound/Dollar exchange rate dataset using a stochastic volatility model. We show in these examples that the AIS method leads to efficient inference when optimally implemented.

The article is organized as follows. Section 2 reviews the original AIS of Neal (2001). Section 3 presents the main results. Section 4 presents the examples and Section 5 concludes. The technical proofs are in the Appendix.

2 Annealed importance sampling

Let $p(y|\theta)$ be the density of the data y, where θ is a parameter vector belonging to a space $\Theta \subset \mathbb{R}^d$. Let $p(\theta)$ be the prior for θ and $\pi(\theta) \propto p(\theta)p(y|\theta)$ its posterior. We are interested in the case where the likelihood $p(y|\theta)$ is analytically intractable but can be unbiasedly estimated. The primary objective in Bayesian inference is to estimate an integral of the form

$$\mathbb{E}_\pi(\varphi) = \int_\Theta \varphi(\theta)\pi(\theta)d\theta \ , \tag{1}$$

for some π-integrable function φ on Θ. We are also interested in estimating the marginal likelihood

$$p(y) = \int p(\theta)p(y|\theta)d\theta. \tag{2}$$

We now present the AIS procedure of Neal (2001) when the likelihood $p(y|\theta)$ can be evaluated pointwise. Let $\pi_0(\theta)$ be some easily-generated density, such as a t density or the prior. Let a_t, $t=0,1,...,T$, be a sequence of real numbers such that $0=a_0<...<a_T=1$, which we call the annealing schedule. A convenient choice is $a_t=t/T$. AIS constructs the following sequence of interpolation densities $\xi_{a_t}(\theta)$, $t=0,...,T$,

$$\xi_{a_t}(\theta) = \frac{\eta_{a_t}(\theta)}{\int \eta_{a_t}(\theta)d\theta}, \quad \text{with} \quad \eta_{a_t}(\theta) = \pi_0(\theta)^{1-a_t}[p(\theta)p(y|\theta)]^{a_t}.$$

Note that $\xi_{a_0}(\theta)=\pi_0(\theta)$ and $\xi_{a_T}(\theta)$ is the posterior $\pi(\theta)$ of interest. Denote by $K_{\xi_{a_t}}(\theta,\cdot)$ a Markov kernel density conditional on θ with invariant distribution ξ_{a_t}, $t=1,...,T-1$. AIS draws M weighted samples $\{w_i,\theta_i\}_{i=1}^M$ as follows.

AIS algorithm. For $i=1,...,M$

- Generate $\theta^{(1)} \sim \xi_{a_0}(\cdot)$.
- For $t=1,...,T-1$, generate $\theta^{(t+1)} \sim K_{\xi_{a_t}}(\theta^{(t)},\cdot)$.
- Set $\theta_i=\theta^{(T)}$ and compute the unnormalized weight

$$w_i = \frac{\eta_{a_1}(\theta^{(1)})}{\eta_{a_0}(\theta^{(1)})} \times \frac{\eta_{a_2}(\theta^{(2)})}{\eta_{a_1}(\theta^{(2)})} \times ... \times \frac{\eta_{a_T}(\theta^{(T)})}{\eta_{a_{T-1}}(\theta^{(T)})}.$$

Note that the algorithm is parallelizable. Neal (2001) shows that the above algorithm is an IS procedure operating on the extended space Θ^T with the artificial target density proportional to

$$f(\theta^{(1)},...,\theta^{(T)}) = \eta_{a_T}(\theta^{(T)})L_{\xi_{a_{T-1}}}(\theta^{(T)},\theta^{(T-1)})...L_{\xi_{a_1}}(\theta^{(2)},\theta^{(1)}),$$

where

$$L_{\xi_{a_t}}(\theta^{(t+1)},\theta^{(t)}) = \frac{\xi_{a_t}(\theta^{(t)})K_{\xi_{a_t}}(\theta^{(t)},\theta^{(t+1)})}{\xi_{a_t}(\theta^{(t+1)})} = \frac{\eta_{a_t}(\theta^{(t)})K_{\xi_{a_t}}(\theta^{(t)},\theta^{(t+1)})}{\eta_{a_t}(\theta^{(t+1)})}$$

is a backward kernel density, and the proposal density

$$g(\theta^{(1)},...,\theta^{(T)}) = \xi_{a_0}(\theta^{(1)})K_{\xi_{a_1}}(\theta^{(1)},\theta^{(2)})...K_{\xi_{a_{T-1}}}(\theta^{(T-1)},\theta^{(T)}).$$

The original target density $\pi = \xi_{a_T}$ is the last marginal of $f(\theta^{(1)},...,\theta^{(T)})$ because

$$\int L_{\xi_{a_t}}(\theta^{(t+1)},\theta^{(t)})d\theta^{(t)} = 1$$

for all t. This shows that AIS is a valid IS method with the weight $w \propto f/g$. Hence, the weighted samples $\{W_i,\theta_i\}_{i=1}^M$ with $W_i = w_i/\sum_{j=1}^M w_j$ approximate $\pi(\theta)$, i.e., $\sum_{i=1}^M W_i\varphi(\theta_i) \to \mathbb{E}_\pi(\varphi)$ almost surely, for any π-integrable function φ.

The AIS procedure explores the parameter space efficiently, and is useful when the target distribution is multimodal and when choosing an appropriate proposal density is challenging.

3 Annealed importance sampling with an estimated likelihood

This section presents an AISEL algorithm for estimating the integral (1) when the likelihood is analytically intractable but can be estimated unbiasedly. Let $\widehat{p}_N(y|\theta)$ denote an estimator of $p(y|\theta)$ with N the number of particles used to estimate the likelihood. We define a sequence of functions $\widetilde{\eta}_{a_t}(\theta)$, $t=0,...,T$, by

$$\widetilde{\eta}_{a_t}(\theta) = \pi_0(\theta)^{1-a_t}[p(\theta)\widehat{p}_N(y|\theta)]^{a_t}.$$

We propose the following algorithm for generating M weighted samples $\{\widetilde{W}_i,\theta_i\}_{i=1}^M$ which approximate the posterior $\pi(\theta)$; see Section 3.1

Algorithm 1 (AISEL). Let $\pi_0(\theta)$ be some easily-generated density.

1. Generate $\theta_i \sim \pi_0(\theta)$, $i=1,...,M$. Set $\widetilde{W}_i = 1/M$, $i=1,...,M$.

2. For $t=1,...,T$

   (i) Weighting: compute the unnormalized weights

$$\widetilde{w}_i = \widetilde{W}_i \frac{\widetilde{\eta}_{a_t}(\theta_i)}{\widetilde{\eta}_{a_{t-1}}(\theta_i)} = \widetilde{W}_i [\pi_0(\theta_i)]^{a_{t-1}-a_t} [p(\theta_i)\widehat{p}_N(y|\theta_i)]^{a_t-a_{t-1}}, \tag{3}$$

   and set the new normalized weights $\widetilde{W}_i = \widetilde{w}_i / \sum_{j=1}^M \widetilde{w}_j$.

   (ii) Resampling: If $\text{ESS} = 1/\sum_{i=1}^M \widetilde{W}_i^2 < \alpha M$ for some $0<\alpha<1$, e.g. $\alpha=1/2$, then resample from $\{\widetilde{W}_i, \theta_i\}_{i=1}^M$, set $\widetilde{W}_i = 1/M$ and (still) denote the resamples by $\{\widetilde{W}_i, \theta_i\}_{i=1}^M$.

   (iii) Markov move: for each $i=1,...,M$, move the sample θ_i according to a Metropolis-Hastings step as follows. Let $q_t(\theta|\theta^c)$ be a proposal with $\theta^c = \theta_i$ the current state. Generate $\theta^p \sim q_t(\theta|\theta^c)$ and set $\theta_i = \theta^p$ with probability

$$\min\left(1, \frac{\pi_0(\theta^p)^{1-a_t}[p(\theta^p)\widehat{p}_N(y|\theta^p)]^{a_t} q_t(\theta^c|\theta^p)}{\pi_0(\theta^c)^{1-a_t}[p(\theta^c)\widehat{p}_N(y|\theta^c)]^{a_t} q_t(\theta^p|\theta^c)}\right).$$

   Otherwise set $\theta_i = \theta^c$.

If $\widehat{p}_N(y|\theta) = p(y|\theta)$, the above algorithm is a special case of the SMC sampler in Del Moral et al. (2006) for sampling from the sequence of distributions ξ_{a_t}. The AIS algorithm of Neal (2001) is a special case of this algorithm in which no resampling steps are performed. It is widely known in the literature that it is beneficial to incorporate resampling steps (Del Moral et al., 2006). The algorithm is also closely related to the resample-move algorithm of Gilks and Berzuini (2001), except that they perform the resampling step in every iteration t.

3.1 Formal justification

The output of Algorithm 1 is weighted samples $\{\widetilde{W}_i, \theta_i\}_{i=1}^M$. To prove that this algorithm is valid, i.e. $\{\widetilde{W}_i, \theta_i\}_{i=1}^M$ approximate $\pi(\theta)$, we make the following assumption.

Assumption 1. $\mathbb{E}[\widehat{p}_N(y|\theta)] = p(y|\theta)$ for every $\theta \in \Theta$.

Let us write $\widehat{p}_N(y|\theta)$ as $p(y|\theta)e^z$ where $z = \log \widehat{p}_N(y|\theta) - \log p(y|\theta)$ is a random variable whose distribution is governed by the randomness occurring when estimating the likelihood $p(y|\theta)$. Let $g_N(z|\theta)$ be the density of z. Assumption 1 implies that

$$\mathbb{E}(e^z) = \int_{\mathbb{R}} e^z g_N(z|\theta) dz = 1.$$

We define

$$\pi_N(\theta, z) = p(\theta) g_N(z|\theta) p(y|\theta) e^z / p(y) \tag{4}$$

as the joint density of θ and z on the extended space $\widetilde{\Theta} = \Theta \otimes \mathbb{R}$. Then its first marginal is the posterior $\pi(\theta)$ of interest, i.e.

$$\int_{\mathbb{R}} \pi_N(\theta, z) dz = \pi(\theta). \tag{5}$$

We define the following sequence of interpolation densities on $\widetilde{\Theta} = \Theta \otimes \mathbb{R}$

$$\widetilde{\xi}_{a_t}(\theta, z) = \frac{\widetilde{\eta}_{a_t}(\theta, z)}{\int \widetilde{\eta}_{a_t}(\theta, z) d\theta dz}, \tag{6}$$

with

$$\widetilde{\eta}_{a_t}(\theta, z) = \pi_0(\theta)^{1-a_t} [p(\theta) p(y|\theta) e^z]^{a_t} g_N(z|\theta), \quad t = 0, ..., T.$$

Note that $\widetilde{\xi}_{a_T}(\theta,z) = \pi_N(\theta,z)$, which is our new target density defined on $\widetilde{\Theta}$. Algorithm 1 is entirely equivalent to the following procedure.

Algorithm 1'.

1. Generate $(\theta_i,z_i) \sim \widetilde{\xi}_{a_0}(\theta,z) = \pi_0(\theta) g_N(z|\theta)$, i.e. generate $\theta_i \sim \pi_0(\theta)$ then $z_i \sim g_N(z|\theta_i)$, $i = 1,...,M$. Set $\widetilde{W}_i = 1/M$.

2. For $t = 1,...,T$

(i) Weighting: compute the weights

$$\widetilde{w}_i = \widetilde{W}_i \frac{\widetilde{\eta}_{a_t}(\theta_i, z_i)}{\widetilde{\eta}_{a_{t-1}}(\theta_i, z_i)} = \widetilde{W}_i [\pi_0(\theta_i)]^{a_{t-1}-a_t} [p(\theta_i)p(y|\theta_i)e^{z_i}]^{a_t-a_{t-1}}.$$

and set the new normalized weights $\widetilde{W}_i = \widetilde{w}_i / \sum_{j=1}^M \widetilde{w}_j$.

(ii) Resampling: If $\mathrm{ESS} < \alpha M$, then resample from $\{\widetilde{W}_i, \theta_i, z_i\}_{i=1}^M$, set $\widetilde{W}_i = 1/M$ and denote the resamples by $\{\widetilde{W}_i, \theta_i, z_i\}_{i=1}^M$.

(iii) Markov move: for each $i = 1,...,M$, move the sample $(\theta^c, z^c) = (\theta_i, z_i)$ by a Metropolis-Hastings kernel $K_{\widetilde{\xi}_{a_t}}(\cdot,\cdot)$ as follows. Generate a proposal (θ^p, z^p) from the proposal density $\widetilde{q}_t(\theta^p, z^p|\theta^c, z^c) = q_t(\theta^p|\theta^c) g_N(z^p|\theta_p)$. Set $(\theta_i, z_i) = (\theta^p, z^p)$ with probability

$$\begin{aligned} \text{prob} &= \min\left(1, \frac{\widetilde{\eta}_{a_t}(\theta^p, z^p)\widetilde{q}_t(\theta^c, z^c|\theta^p, z^p)}{\widetilde{\eta}_{a_t}(\theta^c, z^c)\widetilde{q}_t(\theta^p, z^p|\theta^c, z^c)}\right) \\ &= \min\left(1, \frac{\pi_0(\theta^p)^{1-a_t}[p(\theta^p)\widehat{p}_N(y|\theta^p)]^{a_t} q_t(\theta^c|\theta^p)}{\pi_0(\theta^c)^{1-a_t}[p(\theta^c)\widehat{p}_N(y|\theta^c)]^{a_t} q_t(\theta^p|\theta^c)}\right). \end{aligned}$$

Phrased differently, generating weighted samples $\{\widetilde{W}_i, \theta_i\}_{i=1}^M$ according to Algorithm 1 is equivalent to generating weighted samples $\{\widetilde{W}_i, \theta_i, z_i\}_{i=1}^M$ according to Algorithm 1'. Algorithm 1' is exactly the SMC sampler of Del Moral et al. (2006) for sampling from the sequence $\widetilde{\xi}_{a_t}(\theta, z)$, $t = 0,...,T$, in which the backward kernel used is the backward kernel (30) in their paper. Therefore, the weighted samples $\{\widetilde{W}_i, \theta_i, z_i\}_{i=1}^M$ produced after the last iteration T approximate $\widetilde{\xi}_{a_T}(\theta, z) = \pi_N(\theta, z)$, i.e.

$$\sum_{i=1}^M \widetilde{W}_i \widetilde{\varphi}(\theta_i, z_i) \xrightarrow{a.s.} \int \widetilde{\varphi}(\theta, z)\pi_N(\theta, z) dz d\theta, \quad M \to \infty,$$

for any π_N-integrable function $\widetilde{\varphi}(\theta, z)$ on $\widetilde{\Theta}$. Given the function $\varphi(\theta)$ in (1), we define the corresponding function $\widetilde{\varphi}$ on $\widetilde{\Theta}$ by $\widetilde{\varphi}(\theta, z) = \varphi(\theta)$, then

$$\widehat{\varphi}_{\text{AISEL}} = \sum_{i=1}^{M} \widetilde{W}_i \varphi(\theta_i) = \sum_{i=1}^{M} \widetilde{W}_i \widetilde{\varphi}(\theta_i, z_i) \xrightarrow{a.s.} \int_{\widetilde{\Theta}} \widetilde{\varphi}(\theta, z) \pi_N(\theta, z) dz d\theta$$
$$= \int_{\Theta} \varphi(\theta) \pi(\theta) d\theta$$
$$= \mathbb{E}_{\pi}(\varphi),$$

as $M \to \infty$. This justifies Algorithm 1. We refer to $\widehat{\varphi}_{\text{AISEL}}$ as the AISEL estimator of $\mathbb{E}_{\pi}(\varphi)$.

Remark 1. It is advisable to perform the Markov move step over a few burn-in iterations so that the samples move closer to the equilibrium distribution.

Remark 2. If φ satisfies the conditions in Theorem 1 of Chopin (2004), then

$$\sqrt{M}\Big(\widehat{\varphi}_{\text{AISEL}} - \mathbb{E}_{\pi}(\varphi)\Big) \xrightarrow{d} \mathcal{N}(0, \sigma^2_{\text{AISEL}}(\varphi)) \text{ as } M \to \infty, \tag{7}$$

with the asymptotic variance $\sigma^2_{\text{AISEL}}(\varphi)$ defined recursively as in Chopin (2004).

Except for the special case in which no resampling steps in Algorithm 1 are performed, the asymptotic variance $\sigma^2_{\text{AISEL}}(\varphi)$, and therefore the variance of $\widehat{\varphi}_{\text{AISEL}}$, does not admit a closed form. A natural and potential technique to estimate $\text{Var}(\widehat{\varphi}_{\text{AISEL}})$ is to run Algorithm 1 in batches independently and in parallel. Then, we have several independent batches of weighted samples $\{\widetilde{W}_i^{(r)}, \theta_i^{(r)}\}_{i=1}^{M_r}$, $r = 1,...,R$ with $\sum_r M_r = M$, and the corresponding R independent estimates $\widehat{\varphi}^{(r)}_{\text{AISEL}}$ of $\mathbb{E}_{\pi}(\varphi)$. The variance of the estimator $\widehat{\varphi}_{\text{AISEL}}$ can be estimated by

$$\widehat{\text{Var}}(\widehat{\varphi}_{\text{AISEL}}) = \frac{1}{R} \sum_{r=1}^{R} (\widehat{\varphi}^{(r)}_{\text{AISEL}} - \overline{\widehat{\varphi}})^2 \text{ with } \overline{\widehat{\varphi}} = \frac{1}{R} \sum_{r=1}^{R} \widehat{\varphi}^{(r)}_{\text{AISEL}}.$$

If no resampling steps are performed, then the particles θ_i are independent and we have a closed form expression for estimating the asymptotic variance of $\widehat{\varphi}_{\text{AISEL}}$

$$\widehat{\sigma^2_{\text{AISEL}}(\varphi)} = M \sum_{i=1}^{M} \big(\varphi(\theta_i) - \widehat{\varphi}_{\text{AISEL}}\big)^2 \widetilde{W}_i^2. \tag{8}$$

It is straightforward to show that this estimate is consistent. However, it is important to perform resampling if necessary.

3.2 Estimating the marginal likelihood

Marginal likelihood (2) is important for model comparison purposes. Except for some trivial cases, computing the marginal likelihood is challenging because of its integral form. Friel and Pettitt (2008) propose a very efficient method, called power posterior method, for estimating the marginal likelihood, that exploits the temparing sampling framework as in AIS. This section extends the power posterior method to the case with latent variables.

We consider for now a *continuous* sequence of interpolation densities (6) as follows

$$\widetilde{\xi}_s(\theta, z) = \frac{\widetilde{\eta}_s(\theta, z)}{\int \widetilde{\eta}_s(\theta, z) d\theta dz}, \quad \text{with} \quad \widetilde{\eta}_s(\theta, z) = \pi_0(\theta)^{1-s}[p(\theta)p(y|\theta)e^z]^s g_N(z|\theta),\ 1 \leq s \leq 1.$$

We have the following result. The proof is in the Appendix.

Proposition 1. *Under Assumption 1, the log of the marginal likelihood* $\log p(y)$ *can be expressed as*

$$\log p(y) = \int_0^1 \mathbb{E}_{(\theta,z)\sim\widetilde{\xi}_s} \left[\log \frac{p(\theta)p(y|\theta)e^z}{\pi_0(\theta)}\right] ds. \tag{9}$$

Let $f(s)$ denote the integrand in the right side of (9). The scalar integral in (9) can be deterministically approximated by

$$\widehat{\log p(y)} = \sum_{t=0}^{T-1} (a_{t+1} - a_t) \frac{f(a_{t+1}) + f(a_t)}{2}, \tag{10}$$

with $\{a_t,\ t = 0,...,T\}$ the annealing schedule as in Section 2. The function $f(a_t)$ can be estimated by

$$\widehat{f}(a_t) = \sum_{i=1}^{M} \widetilde{W}_i^{(t)} \log \frac{p(\theta_i^{(t)})\widehat{p}_N(y|\theta_i^{(t)})}{\pi_0(\theta_i^{(t)})} \tag{11}$$

with $\{\widetilde{W}_i^{(t)}, \theta_i^{(t)}\}_{i=1}^M$ the output of Algorithm 1 after iteration t.

This approach of estimating the marginal likelihood fits naturally to the AISEL procedure

and is straightforward to implement. Note that the values $\widehat{p}_N(y|\theta_i^{(t)})$ in (11) can be used for the weighting step in iteration $t+1$ of Algorithm 1, hence no extra computation is needed except calculation in (10) and (11).

3.3 The effect of estimating the likelihood

The efficiency of an IS procedure with proposal density g and weights w_i is often measured by the effective sample size defined by (see, e.g. Neal (2001) and Liu (2001, Chapter 2))

$$\text{ESS} = \frac{M}{1+\text{Var}_g(w_i/\mathbb{E}_g[w_i])} = \frac{M}{1+\text{CV}_g(w_i)},$$

where $\text{CV}_g(w_i) = \text{Var}_g(w_i)/(\mathbb{E}_g[w_i])^2$ is often called the coefficient of variation of the unnormalized weights w_i. The bigger the ESS the more efficient the IS procedure. This section investigates how much the ESS is reduced when working with an estimated likelihood.

We consider the case of the original AIS procedure, i.e. Algorithm 1 without the resampling step, and work with the notation in Algorithm 1'. Write $x=(\theta,z)$ and $x^{(t)}=(\theta^{(t)},z^{(t)})$. Without the resampling step, Algorithm 1' can be written as

- Generate $x^{(1)}=(\theta^{(1)},z^{(1)}) \sim \widetilde{\xi}_{a_0}(\theta,z) = \pi_0(\theta)g_N(z|\theta)$.
- For $t=1,...,T-1$, generate $x^{(t+1)}=(\theta^{(t+1)},z^{(t+1)})$ from the Markov kernel $K_{\widetilde{\xi}_{a_t}}(x^{(t)},\cdot)$.
- Set $(\theta_i,z_i)=(\theta^{(T)},z^{(T)})$ and compute the corresponding unnormalized weight

$$\widetilde{w}_i = \frac{\widetilde{\eta}_{a_1}(x^{(1)})}{\widetilde{\eta}_{a_0}(x^{(1)})} \times \frac{\widetilde{\eta}_{a_2}(x^{(2)})}{\widetilde{\eta}_{a_1}(x^{(2)})} \times ... \times \frac{\widetilde{\eta}_{a_T}(x^{(T)})}{\widetilde{\eta}_{a_{T-1}}(x^{(T)})}.$$

Denote by ESS_{AIS} and $\text{ESS}_{\text{AISEL}}$ the effective sample sizes of the AIS procedures when the likelihood is given and when it is estimated, respectively.

We make the following assumption which is satisfied in almost cases.

Assumption 2. There exists a function $\lambda(\theta)$ such that for each $\theta \in \Theta$, $\lambda^2(\theta) < \infty$ and

$$\sqrt{N}\left(\widehat{p}_N(y|\theta) - p(y|\theta)\right) \xrightarrow{d} \mathcal{N}(0, \lambda^2(\theta)) \text{ as } N \to \infty.$$

The following lemma follows immediately from Assumption 2 using the second order δ-method.

Lemma 1. *Let $\gamma^2(\theta) = \lambda^2(\theta)/p(y|\theta)^2$, and suppose that Assumption 2 holds. Let $z \sim g_N(z|\theta)$. Then*

$$\sqrt{N}\left(\frac{z + \frac{\gamma^2(\theta)}{2N}}{\gamma(\theta)}\right) \xrightarrow{d} \mathcal{N}(0,1) \text{ as } N \to \infty.$$

Following Pitt et al. (2012), we make the following further assumptions.

Assumption 3. (i) The density $g_N(z|\theta)$ of z is $\mathcal{N}(-\frac{\gamma^2(\theta)}{2N}, \frac{\gamma^2(\theta)}{N})$. (ii) For a given $\sigma^2 > 0$, let N be a function of θ and σ^2 such that $\mathrm{Var}(z) \equiv \sigma^2$, i.e. $N = N_{\sigma^2}(\theta) = \gamma^2(\theta)/\sigma^2$.

Assumption 3(i) is justified by Lemma 1. Assumption 3(ii) keeps the variance $\mathrm{Var}(z)$ constant across different values of θ, thus making it easy to associate the ESS with σ. Under Assumption 3, the density $g_N(z|\theta)$ depends only on σ and is denoted by $g(z|\sigma)$.

Assumption 4. $K_{\xi_{a_t}}(\theta, \cdot) = \xi_{a_t}(\cdot)$ and $K_{\widetilde{\xi}_{a_t}}(x, \cdot) = \widetilde{\xi}_{a_t}(\cdot)$.

As in Neal (2001), this assumption separates out the effect related to Markov chain convergence and allows us to study the effect of estimating the likelihood on the sequential sampling scheme.

Theorem 1. *Suppose that Assumptions 1-4 hold. Then*

$$\frac{\mathrm{ESS}_{\mathrm{AISEL}}}{\mathrm{ESS}_{\mathrm{AIS}}} = \exp\left(-\tau\sigma^2\right), \tag{12}$$

with $\tau = \sum_{t=1}^{T}(a_t - a_{t-1})(2a_t - 1) > 0$ for any sequence $0 = a_0 < a_1 < \ldots < a_T = 1$.

The theorem, whose proof is in the Appendix, shows that the efficiency is reduced by the factor $\exp(\tau\sigma^2)$ when working with an estimated likelihood. If $a_t = 1/T$, then $\tau = 1/T$ for all t, and the theorem shows that increasing T and thus making the $\widetilde{\xi}_{a_t}$ closer to each other helps improve efficiency.

In the IS^2 approach, Tran et al. (2013) show that $\mathrm{ESS}_{\mathrm{IS}^2}/\mathrm{ESS}_{\mathrm{IS}} = \exp(-\sigma^2)$, where $\mathrm{ESS}_{\mathrm{IS}^2}$ and $\mathrm{ESS}_{\mathrm{IS}}$ are the effective sample sizes of IS when the likelihood is estimated and given, respectively. Similarly in particle MCMC, Pitt et al. (2012) show that efficiency is reduced by a factor of approximately $\exp(-\sigma^2)$ when working with an estimated likelihood. We can see that the accuracy of likelihood estimation is less important in AISEL than in IS^2 and particle MCMC, because the factor $\exp(-\tau\sigma^2)$ can be made small when τ is decreased. This means that AISEL can be more robust than IS^2 and particle MCMC in cases where we only have a rough estimate of the likelihood, or it is expensive to obtain an accurate estimate of the likelihood.

3.4 Practical guidelines on selecting the number of particles

This section studies how to select the number of particles N optimally. A large number of particles N results in a precise likelihood estimate, and therefore an accurate estimate of $\mathbb{E}_\pi(\varphi)$, but at a greater computational cost. A small N leads to a large variance of the likelihood estimator, so we need a larger number of importance samples M in order to obtain the desired accuracy of the AISEL estimator. In either case, the computation is expensive. It is important to select an optimal value of N that minimizes the computational cost.

The time to compute the likelihood estimate $\widehat{p}_N(y|\theta)$ can be written as $\tau_0 + N(\theta)\tau_1$ where $\tau_0 \geq 0$ and $\tau_1 > 0$ (Tran et al., 2013). For example, if $\widehat{p}_N(y|\theta)$ is estimated by IS, then τ_0 is the overhead cost spent on estimating the proposal density and τ_1 is the computing time used to generate each sample and compute the weight. Note that under Assumption 3, N depends on θ as $N = N_{\sigma^2}(\theta) = \gamma^2(\theta)/\sigma^2$.

The variance of the AISEL estimator is approximated as

$$\mathrm{Var}(\widehat{\varphi}_{\mathrm{AISEL}}) \approx \frac{\mathrm{Var}_\pi(\varphi)}{\mathrm{ESS}_{\mathrm{AISEL}}},$$

with $\mathrm{Var}_\pi(\varphi)=\mathbb{E}_\pi(\varphi-\mathbb{E}_\pi(\varphi))^2$. From (12),

$$\mathrm{Var}(\widehat{\varphi}_{\mathrm{AISEL}}) \approx \frac{1}{M}\mathrm{Var}_\pi(\varphi)(1+\mathrm{CV}_g(w_i))\exp(\tau\sigma^2). \tag{13}$$

Let P^* be a prespecified precision. Then we need approximately

$$M(P^*) = \frac{1}{P^*}\mathrm{Var}_\pi(\varphi)(1+\mathrm{CV}_g(w_i))\exp(\tau\sigma^2)$$

particles in order to have that precision. The required computing time to run AISEL is

$$\sum_{t=1}^{T}\sum_{i=1}^{M}(N(\theta_i^{(t)})\tau_1+\tau_0) \approx TM\left(\frac{\bar{\gamma}^2}{\sigma^2}\tau_1+\tau_0\right) = \frac{T}{P^*}\mathrm{Var}_\pi(\varphi)(1+\mathrm{CV}_g(w_i))\exp(\tau\sigma^2)\left(\frac{\bar{\gamma}^2}{\sigma^2}\tau_1+\tau_0\right), \tag{14}$$

in which

$$\frac{1}{TM}\sum_{t=1}^{T}\sum_{i=1}^{M}\gamma^2(\theta_i^{(t)}) \longrightarrow \bar{\gamma}^2 = \frac{1}{T}\sum_{t=1}^{T}\mathbb{E}_{\xi_{a_t}}[\gamma^2(\theta)], \quad M\to\infty.$$

Therefore

$$\mathrm{CT}^*(\sigma^2) = \exp(\tau\sigma^2)\times\left(\frac{\bar{\gamma}^2}{\sigma^2}\tau_1+\tau_0\right) \tag{15}$$

characterizes the computing as a function of σ^2, which is minimized at

$$\sigma^2_{\mathrm{opt}} = \begin{cases} \frac{\sqrt{(\bar{\gamma}^2\tau\tau_1)^2+4\bar{\gamma}^2\tau\tau_0\tau_1}-\bar{\gamma}^2\tau\tau_1}{2\tau\tau_0}, & \tau_0>0 \\ 1/\tau, & \tau_0=0. \end{cases} \tag{16}$$

The optimal number of particles N is such that

$$\mathrm{Var}_{N,\theta}(z) = \mathrm{Var}(\log\widehat{p}_N(y|\theta)) = \sigma^2_{\mathrm{opt}}.$$

Let $\widehat{\mathrm{Var}}_{N,\theta}(z)$ be an estimate of $\mathrm{Var}_{N,\theta}(z)$, which can be obtained by using, e.g., the delta method or the jackknife. See Tran et al. (2013) for more details. We suggest the following practical guidelines for tuning the optimal number of particles N. Note that N generally

depends on θ but this dependence is suppressed for notational simplicity.

The case $\tau_0 = 0$. From (15), $\sigma^2_{\text{opt}} = 1/\tau$. It is necessary to tune N such that $\widehat{\text{Var}}_{N,\theta}(z) = 1/\tau$. A simple strategy is to start with some small N and increase it if $\widehat{\text{Var}}_{N,\theta}(z) > 1/\tau$.

The case $\tau_0 > 0$. First, we need to estimate $\bar{\gamma}^2$. Let $\{\theta_1,...,\theta_J\}$ be a few initial draws from the initial density $\xi_0(\theta)$. Then, start with some large N_0, $\bar{\gamma}^2$ can be initially estimated by

$$\hat{\bar{\gamma}}^2 = \frac{1}{J}\sum_{j=1}^{J} \hat{\gamma}^2(\theta_j) = \frac{N_0}{J}\sum_{j=1}^{J} \widehat{\text{Var}}_{N_0,\theta_j}(z), \tag{17}$$

as $\widehat{\text{Var}}_{N_0,\theta_j}(z) = \hat{\gamma}^2(\theta_j)/N_0$. By substituting this estimate of $\bar{\gamma}^2$ into (16) we obtain an estimate $\widehat{\sigma}^2_{\text{opt}}$ of σ^2_{opt}. We now can start Algorithm 1 and update $\hat{\bar{\gamma}}^2$ (and therefore $\widehat{\sigma}^2_{\text{opt}}$) as we go. For each draw of θ, we start with some small N and increase N if $\widehat{\text{Var}}_{N,\theta}(z) > \widehat{\sigma}^2_{\text{opt}}$.

Time normalized variance. In the examples in Section 4 we use the time normalized variance (TNV) as a measure of efficiency (Tran et al., 2013) of a sampling procedure. The TNV of the AISEL estimator $\widehat{\varphi}_{\text{AISEL}}$ is defined as

$$\text{TNV}(M, N) = \text{Var}(\widehat{\varphi}_{\text{AISEL}}) \times \tau(M, N), \tag{18}$$

where $\tau(M,N)$ is the total CPU time used to run the AISEL procedure with M importance samples and N particles. From (13) and (14),

$$\text{TNV}(M, N) \approx T\text{Var}_\pi(\varphi)(1 + \text{CV}_g(w_i)) \exp(\tau\sigma^2) \left(\frac{\bar{\gamma}^2}{\sigma^2}\tau_1 + \tau_0\right)$$

and is propotional to $\text{CT}^*(\sigma^2)$.

Remark 3. Letting the optimal $N_{\text{opt}} = N_{\text{opt}}(\theta)$ depend on θ is theoretically interesting but might in some cases be ultimately inefficient. The reason is that extra computing time is needed to tune N for each θ. A simple strategy is to make approximation that $\gamma^2(\theta) \approx \bar{\gamma}^2$.

Then, the optimal number of particles, which is constant across θ, is determined as

$$N_{\text{opt}} = \begin{cases} \frac{2\tau\tau_0}{\sqrt{(\tau\tau_1)^2+4\tau\tau_0\tau_1/\bar{\gamma}^2}-\bar{\gamma}^2\tau\tau_1}, & \tau_0 > 0 \\ \tau\bar{\gamma}^2, & \tau_0 = 0. \end{cases} \tag{19}$$

We follow this strategy in the examples below.

4 Examples

4.1 A simulation example

We generate a dataset from a mixed logistic regression model

$$P(y_{ij} = 1|\beta, \eta_i) = \frac{\exp(\beta_0 + x'_{ij}\beta + \eta_{i0} + z_{ij}\eta_{i1})}{1 + \exp(\beta_0 + x'_{ij}\beta + \eta_{i0} + z_{ij}\eta_{i1})}, \;\; j = 1, ..., n_i, \;\; i = 1, ..., m$$

in which the random effects $\eta_i = (\eta_{i0},\eta_{i1})' \sim N(0,\Sigma)$ and $\Sigma = \text{diag}(\sigma_1^2,\sigma_2^2)$. Covariates are generated from the uniform distribution $U(0,1)$. We set $\beta_0 = -3$, $\beta = (2, \; -2, \; 2)'$, $\sigma_1^2 = 2$, $\sigma_2^2 = 1$, $n_i = 10$ and $m = 50$. We use a normal prior $N(0,100I)$ for β and $p(\sigma_k^2) \propto 1/\sigma_k^2$, $k = 1,2$. Algorithm 1 is used to estimate the posterior mean of the parameters $\theta = (\beta_0,\beta,\sigma_1^2,\sigma_2^2)$. The sequence a_t is set as $1/T$ with $T = 10$. The initial distribution π_0 is a multivariate t with mean (0, 0, 0, 0, 1, 2), scale matrix $3I_6$ and degrees of freedoms 10.

We first run Algorithm 1 to generate $M = 100$ importance samples with $N = 10$, and obtain $\tau_0 = 7.2 \times 10^{-3}$,$\tau_1 = 5.9 \times 10^{-4}$ and $\bar{\gamma}^2 = 17.7$. This gives $\sigma^2_{\text{opt}} = 2.6$ and the optimal number of particles (as constant across θ) $N_{\text{opt}} = 7$. We then run Algorithm 1 to generate $M = 5000$ importance samples for five values of N, $N = 1$,7,10,20 and $N = 50$. In order to be able to estimate the variance of the estimator, Algorithm 1 is run in parallel as described in Remark 2 with $R = 20$ batches.

Figure 1 plots the time normalized variance in (18) versus N. The TNV is averaged over the posterior mean estimates of the four parameters β_0, β, σ_1^2 and σ_2^2. The TNV appears to be

minimized at $N=10$, close to the theoretical optimal value $N=7$. The results suggests that the TNV is weakly sensitive around the optimal value of N. The efficiency decreases linearly when N is higher than the optimal value, whereas the efficiency can deteriorate exponentially when N is below the optimal. This phenomenon is also observed in the IS^2 method (Tran et al., 2013). In practice, it is therefore advisable to use for N a value which is slightly bigger than N_{opt}.

Table 1 reports the estimate of the posterior mean when using $N=10$ particles.

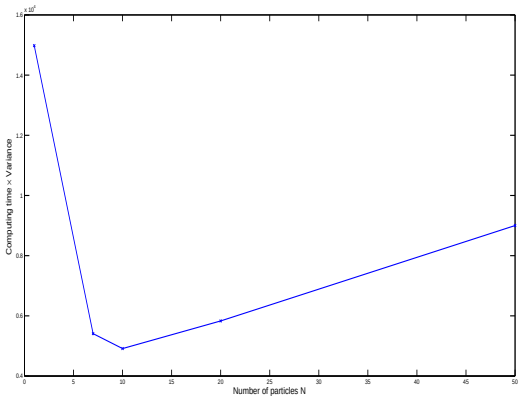

Figure 1: Plot of the time normalized variance vs N

	True	Mean	Std. Dev
β_0	-3	-3.08	0.40
β_1	2	1.97	0.08
β_2	-2	-1.99	0.05
β_3	2	2.02	0.08
σ_1^2	2	2.15	0.68
σ_2^2	1	0.62	0.07

Table 1: Estimate of the posterior mean and the standard deviation

4.2 Real Data Example: The SV model

We analyze the Pound/Dollar data set (Kim et al. (1998)) using both the standard SV model and the SV model with leverage effect. For the standard SV model the measurement equation,

for the t^{th} observation y_t, is given by

$$y_t = \exp\left(\frac{h_t}{2}\right)\varepsilon_t;\ t = 1, 2, \ldots, n, \tag{20}$$

where ε_t follows the standard normal distribution, and h_t is a latent variable defined for $t=1,2,...,n-1$, as

$$h_{t+1} = \mu\left(1-\phi\right) + \phi h_t + \sigma_\eta \eta_t, \tag{21}$$

where μ is the unconditional mean, ϕ is the level of persistence, σ_η is the scaling parameter, for the latent process and η_t follows the standard normal distribution. The SV model given in (20) and (21) is completed by the initial state, given as

$$h_1 \sim N\left(m_1, v_1\right). \tag{22}$$

For the analysis we assume, *a priori*, that μ is normally distributed, such that $\mu \sim N(0,100)$, we assume that the log of variance, h_t, is generated by a stationary process with positive autocorrelation, and where the persistence parameter, ϕ, follows a Beta distribution where $\phi \sim Be(15,1.5)$ and we assume σ_η has an inverted gamma prior $IG(10,0.1)$.

To implement Algorithm 1, we use the bootstrap particle filter of Gordon et al. (1993) to obtain an unbiased estimate of the log-likelihood. For this specific implementation $\tau_0 = 0$, for which (16) implies that the variance of the estimated log-likelihood should be $\frac{1}{\tau}$ for the annealed IS scheme that we implement. In the analysis of the SV model we set $T = 15$. To estimate the optimal number particles, we compute the log-likelihood one thousand times, evaluated at parameter values that are typical for the SV model. Specifically, we set $\mu = -0.6$, $\phi = 0.98$ and we set $\sigma_\eta = 0.16$. This leads us to the conclusion that 24 is the optimal number of particles for this model and data set. Algorithm 1 also requires us to set M, which we set to $M = 1000$.

We also need to specify step (iii) of Algorithm 1, which is the Markov move step. Specifically, we employ 5 random walk Metropolis Hastings (RWMH) steps for each Markov move;

see Robert and Casella (1999) for further details on the RWMH algorithm. Similarly, to Chopin (2002) (in the sequential Monte Carlo context), we take the covariance of the RWMH algorithm as the covariance of the current set of (annealed IS) particles. The covariance is scaled by the parameter α, where α is adjusted at each move step, based on the acceptance rate of the previous move step. Specifically, at each move step we update the scale parameter, α, such that

$$\alpha \leftarrow MF \times \alpha, \tag{23}$$

where MF denotes a multiplication factor that scales α from the previous period. The multiplication factor is determined by the acceptance rate (AR) from the previous Markov move step.

Range AR	[0,0.01)	[0.01,0.1)	[0.1,0.15)	[0.15,0.2)	[0.2,0.23)
MF	0.2	0.5	0.7	0.9	0.99
Range AR	[0.23,0.25)	[0.25,0.5)	[0.5,0.85)	[0.85,0.99)	[0.99,1.0)
MF	1	1/0.97	1/ 0.8	1/0.7	1/0.5

Table 2: reports the value of the multiplication factor, MF, given the acceptance rate of the previous Markov move step

We report in Table 2 the multiplication factor used in (23), given the value of the acceptance rate from the previous period. We find adapting in this fashion works well in all the examples we have considered so far, including the ones we consider in this paper.

The output for Figure 2 is produced by running Algorithm 1, using the method described in Remark 3 with $R = 200$ batches. For this part of the analysis here we set $a_t = t/n$. As in the simulated example, the theoretically optimal number of simulated particles, is close the what is empirically observed as optimal. We also observe that the penalty for using too few particles, with respect to the time corrected measure of accuracy, can be much greater than using too many particles.

We use the annealed IS algorithm for estimation of the SV model on the Pound/Dollar data set. Here we set $a_t = (t/n)^3$. This ensures that we move away from the prior very

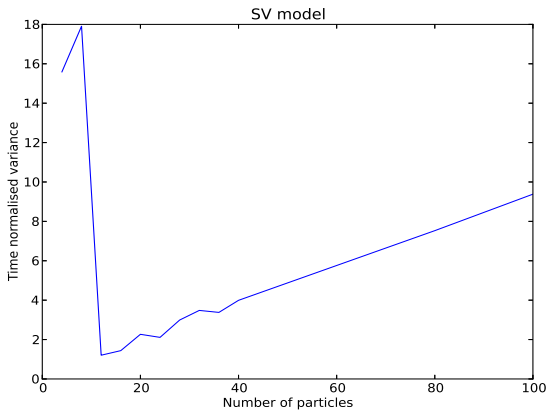

Figure 2: Plot of the time normalized variance vs N for the SV model

	Mean	Std. Dev
μ	-0.66	0.40
ϕ	0.98	0.02
σ_η	0.17	0.04
log ML	-19	

Table 3: Reports output from the analysis of the SV model on the Pound/Dollar data set.

slowly initially, and ensures that the effective sample size is high as we move over the entire extended state space. This is important when estimating the marginal likelihood as we require accurate estimates of the expected value of the log-likelihood across the entire temperature range, and not just at the end of the estimation process, which is the case for parameter estimates. The results are reported in Table 3. We report both the posterior mean and standard deviation for each of the parameters. The time taken for estimation is 176 seconds, using the Julia programming language (Bezanson et al. (2012)), on a Core i7 Linux box, with 16 Gigabytes of RAM. Note the code has not been parallelized, so even better performance could be achieved, with a more highly optimized implementation.

Unlike the standard SV model, in (20), (21) and (22), in which the measurement and state disturbances are independent, the SV model with leverage effect allows for correlation

between ε_t and η_t; see Omori et al. (2007), for further details. Specifically, it is assumed that

$$\begin{pmatrix} \varepsilon_t \\ \eta_t \end{pmatrix} \sim N\left(0, \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}\right),$$

where $|\rho|<1$. As for the standard SV model we use the bootstrap particle filter to obtain an estimate of the log-likelihood. We set the number of particles to 20, which corresponds to our estimate of the optimal number of particles. The other algorithmic parameters remain the same as for the standard SV model.

	Mean	Std. Dev
μ	-0.71	0.26
ϕ	0.98	0.02
σ_η	0.17	0.04
ρ	-0.04	0.12
log ML	-73	

Table 4: Reports output from the analysis of the SV model with leverage on the Pound/Dollar data set.

Output from the analysis of the Pound/Dollar data set using the SV model with leverage is reported in Table 4. The analysis took 172 seconds on using the Julia programming language (Bezanson et al. (2012)), on a Core i7 Linux box, with 16 Gigabytes of RAM. There is little evidence of leverage effect, for the Pound/Dollar data set, based on this analysis. In particular, the log marginal likelihood strongly favours the standard SV model, which isn't surprising given that the estimate of ρ is close to zero.

5 Conclusions

We have presented the annealed IS algorithm for Bayesian inference in models with latent variables. The proposed AISEL method can be considered as a supplement to existing Monte Carlo methods for latent variable models, including particle MCMC (Beaumont, 2003; Andrieu and Roberts, 2009; Andrieu et al., 2010), SMC^2 (Chopin et al., 2013) and IS^2 (Tran et al., 2013). The theory and methodology presented in this paper are useful for

Bayesian inference in latent variable models where the posterior distribution is multimodal and choosing an appropriate proposal density is challenging. An estimate of the log marginal likelihood is obtained as a byproduct of the estimation procedure.

Acknowledgment

The research was partially supported by Australian Research Council grant DP0667069.

Appendix

Proof of Theorem 1. As in Section 2, we can see that the AISEL algorithm is an IS procedure operating on the extended space $(\Theta \otimes \mathbb{R})^T$ with the proposal density

$$\widetilde{g}(x^{(1)}, ..., x^{(T)}) = \widetilde{\xi}_{a_0}(x^{(1)}) K_{\widetilde{\xi}_{a_1}}(x^{(1)}, x^{(2)}) ... K_{\widetilde{\xi}_{a_{T-1}}}(x^{(T-1)}, x^{(T)}).$$

Note that $\widetilde{\eta}_{a_t}(\theta,z) = \eta_{a_t}(\theta)\exp(a_t z) g_N(z|\theta)$. We have

$$\begin{aligned} \log(\widetilde{w}_i) &= \sum_{t=1}^{T} \left(\log \widetilde{\eta}_{a_t}(x^{(t)}) - \log \widetilde{\eta}_{a_{t-1}}(x^{(t)}) \right) \\ &= \sum_{t=1}^{T} \left(\log \eta_{a_t}(\theta^{(t)}) - \log \eta_{a_{t-1}}(\theta^{(t)}) \right) + \sum_{t=1}^{T} (a_t - a_{t-1}) z_t \\ &= \log(w_i) + \sum_{t=1}^{T} (a_t - a_{t-1}) z_t. \end{aligned} \tag{24}$$

Denote $\widetilde{\theta} = (\theta^{(1)},...,\theta^{(T)})$ and $\widetilde{z} = (z^{(1)},...,z^{(T)})$. Under Assumption 4,

$$g_\theta(\widetilde{\theta}) = g_\theta(\theta^{(1)}, ..., \theta^{(T)}) = \xi_{a_0}(\theta^{(1)}) \xi_{a_1}(\theta^{(2)}) ... \xi_{a_{T-1}}(\theta^{(T)}).$$

and

$$\widetilde{g}(x^{(1)}, ..., x^{(T)}) = \widetilde{\xi}_{a_0}(x^{(1)}) \widetilde{\xi}_{a_1}(x^{(2)}) ... \widetilde{\xi}_{a_{T-1}}(x^{(T)}).$$

By Assumption 3,

$$\widetilde{\xi}_{a_t}(\theta, z) = \xi_{a_t}(\theta) g_t(z) \quad \text{with} \quad g_t(z) = \frac{1}{C_t} e^{a_t z} g(z|\sigma), \quad C_t = \exp(-\frac{1}{2} a_t \sigma^2 + \frac{1}{2} a_t^2 \sigma^2).$$

Hence

$$\widetilde{g}(x^{(1)}, ..., x^{(T)}) = \prod_{t=1}^{T} \xi_{a_{t-1}}(\theta^{(t)}) \prod_{t=1}^{T} g_{t-1}(z^{(t)}) = g_\theta(\widetilde{\theta}) g_z(\widetilde{z})$$

with

$$g_z(\widetilde{z}) = \prod_{t=1}^{T} g_{t-1}(z^{(t)}).$$

From (24), we have

$$\begin{aligned}
\mathbb{E}_{\widetilde{g}}[\widetilde{w}_i^2] &= \int w_i^2(\widetilde{\theta}) \prod_{t=1}^{T} e^{2(a_t - a_{t-1}) z^{(t)}} g_z(\widetilde{z}) g_\theta(\widetilde{\theta}) d\widetilde{z} d\widetilde{\theta} \\
&= \int w_i^2(\widetilde{\theta}) g_\theta(\widetilde{\theta}) d\widetilde{\theta} \times \prod_{t=1}^{T} \int e^{2(a_t - a_{t-1}) z^{(t)}} g_{t-1}(z^{(t)}) dz^{(t)} \\
&= \mathbb{E}_g[w_i^2] \prod_{t=1}^{T} \frac{1}{C_{t-1}} \exp\left(-\frac{1}{2}(2a_t - a_{t-1})\sigma^2 + \frac{1}{2}(2a_t - a_{t-1})^2 \sigma^2 \right) \\
&= \mathbb{E}_g[w_i^2] \exp\left(\sigma^2 \sum_{t=1}^{T} (a_t - a_{t-1})(2a_t - 1) \right) \\
&= \mathbb{E}_g[w_i^2] \exp\left(\tau \sigma^2 \right)
\end{aligned}$$

with $\tau = \sum_{t=1}^{T} (a_t - a_{t-1})(2a_t - 1)$. Note that $\mathbb{E}_{\widetilde{g}}[\widetilde{w}_i] = \mathbb{E}_g[w_i]$ by the unbiasedness property of IS. Hence,

$$\begin{aligned}
1 + \mathrm{CV}_{\widetilde{g}}(\widetilde{w}_i) &= 1 + \frac{\mathrm{Var}_{\widetilde{g}}(\widetilde{w}_i)}{(\mathbb{E}_{\widetilde{g}}[\widetilde{w}_i])^2} = \frac{\mathbb{E}_{\widetilde{g}}[\widetilde{w}_i^2]}{(\mathbb{E}_{\widetilde{g}}[\widetilde{w}_i])^2} = \frac{\mathbb{E}_g[w_i^2]}{(\mathbb{E}_g[w_i])^2} \exp(\tau \sigma^2) \\
&= (1 + \mathrm{CV}_g(w_i)) \exp(\tau \sigma^2). \qquad (25)
\end{aligned}$$

So

$$\mathrm{ESS}_{\mathrm{AISEL}} = \frac{M}{1 + \mathrm{CV}_{\widetilde{g}}(\widetilde{w}_i)} = \frac{M}{1 + \mathrm{CV}_g(w_i)} \exp(-\tau \sigma^2) = \exp(-\tau \sigma^2) \mathrm{ESS}_{\mathrm{AIS}}. \qquad (26)$$

Note that

$$\tau = \sum_{t=1}^{T}(2a_t^2 - 2a_t a_{t-1}) - \sum_{t=1}^{T}(a_t - a_{t-1}) > \sum_{t=1}^{T}(2a_t^2 - a_t^2 - a_{t-1}^2) - 1 = 0.$$

□

Proof of Proposition 1. The proof follows Friel and Pettitt (2008) who prove the result in the case $\pi_0(\theta)=p(\theta)$ and the likelihood $p(y|\theta)$ is analytically available. Let

$$\zeta(s) = \int \widetilde{\eta}_s(\theta, z)d\theta dz,\ 0 \leq s \leq 1.$$

Then, $\zeta(0)=1$ and by Assumption 1, $\zeta(1)=p(y)$. Note that $\widetilde{\xi}_s(\theta,z)=\widetilde{\eta}_s(\theta,z)/\zeta(s)$.

$$\begin{aligned}
\frac{d\zeta(s)}{ds} &= \frac{d}{ds}\int \pi_0(\theta)\left(\frac{p(\theta)p(y|\theta)e^z}{\pi_0(\theta)}\right)^s g_N(z|\theta)d\theta dz \\
&= \int \pi_0(\theta)\left(\frac{p(\theta)p(y|\theta)e^z}{\pi_0(\theta)}\right)^s \left[\log\frac{p(\theta)p(y|\theta)e^z}{\pi_0(\theta)}\right] g_N(z|\theta)d\theta dz \\
&= \int \widetilde{\eta}_s(\theta, z)\left[\log\frac{p(\theta)p(y|\theta)e^z}{\pi_0(\theta)}\right] d\theta dz.
\end{aligned}$$

Hence,

$$\begin{aligned}
\frac{d\log\zeta(s)}{ds} &= \frac{1}{\zeta(s)}\frac{d\zeta(s)}{ds} \\
&= \int \widetilde{\xi}_s(\theta, z)\left[\log\frac{p(\theta)p(y|\theta)e^z}{\pi_0(\theta)}\right] d\theta dz \\
&= \mathbb{E}_{(\theta,z)\sim\widetilde{\xi}_s}\left[\log\frac{p(\theta)p(y|\theta)e^z}{\pi_0(\theta)}\right].
\end{aligned}$$

So

$$\int_0^1 \mathbb{E}_{(\theta,z)\sim\widetilde{\xi}_s}\left[\log\frac{p(\theta)p(y|\theta)e^z}{\pi_0(\theta)}\right] ds = \log\zeta(1) - \log\zeta(0) = \log p(y).$$

□

References

Andrieu, C., Doucet, A., and Holenstein, R. (2010). Particle Markov chain Monte Carlo methods. *Journal of the Royal Statistical Society, Series B*, 72:1–33.

Andrieu, C. and Roberts, G. (2009). The pseudo-marginal approach for efficient Monte Carlo computations. *The Annals of Statistics*, 37:697–725.

Beaumont, M. A. (2003). Estimation of population growth or decline in genetically monitored populations. *Genetics*, 164:1139–1160.

Bezanson, J., Karpinski, S., Shah, V. B., and Edelman, A. (2012). Julia: A fast dynamic language for technical computing. *CoRR*, abs/1209.5145.

Chopin, N. (2002). A sequential particle method for static models. *Biometrika*, 89:539–551.

Chopin, N. (2004). Central limit theorem for sequential Monte Carlo methods and its application to Bayesian inference. *The Annals of Statistics*, 32(6):2385–2411.

Chopin, N., Jacob, P. E., and Papaspiliopoulos, O. (2013). SMCˆ2: an efficient algorithm for sequential analysis of state-space models. *Journal of the Royal Statistical Society B*, 75(3):397–426.

Del Moral, P., Doucet, A., and Jasra, A. (2006). Sequential Monte Carlo samplers. *Journal of the Royal Statistical Society, Series B*, 68:411–436.

Durbin, J. and Koopman, S. J. (2001). *Time Series Analysis by State Space Methods*. Oxford University Press.

Filippone, M. (2013). Bayesian inference for Gaussian process classifiers with annealing and exact-approximate MCMC. Technical report, University of Glasgow.

Fitzmaurice, G. M., Laird, N. M., and Ware, J. H. (2011). *Applied Longitudinal Analysis*. John Wiley & Sons, Ltd, New Jersey, 2nd edition.

Friel, N. and Pettitt, A. N. (2008). Marginal likelihood estimation via power posteriors. *J. R. Statist. Soc. B*, 70:589–607.

Gilks, W. R. and Berzuini, C. (2001). Following a moving target - Monte Carlo inference for dynamic Bayesian models. *Journal of the Royal Statistical Society, Series B*, 63.

Gordon, N. J., Salmond, D. J., and Smith, A. F. M. (1993). A novel approach to non-linear and non-Gaussian Bayesian state estimation. *Radar and Signal Processing, IEE Proceedings*, 140:107–113.

Kim, S., Shephard, N., and Chib, S. (1998). Stochastic volatility: Likelihood inference and comparison with ARCH models. *Review of Economic Studies*, 65(3):361–393.

Liu, J. S. (2001). *Monte Carlo Strategies in Scientific Computing*. Springer-Verlag, New York.

Neal, R. (2001). Annealed importance sampling. *Statistics and Computing*, 11:125–139.

Omori, Y., Chib, S., Shephard, N., and Nakajima, J. (2007). Stochastic volatility with leverage: Fast and efficient likelihood inference. *Journal of Econometrics*, 140:425–449.

Pitt, M. K., Silva, R. S., Giordani, P., and Kohn, R. (2012). On some properties of Markov chain Monte Carlo simulation methods based on the particle filter. *Journal of Econometrics*, 171(2):134–151.

Rasmussen, C. E. and Williams, C. (2006). *Gaussian Processes for Machine Learning*. MIT Press.

Robert, C. P. and Casella, G. (1999). *Monte Carlo Statistical Methods*. Springer–Verlag, New York.

Tran, M.-N., Scharth, M., Pitt, M. K., and Kohn, R. (2013). Importance sampling squared for Bayesian inference in latent variable models. Technical report, University of New South Wales.